Interim Progress Report
Sandia Red Team Assessment of Select DARPA ITO Projects

Executive Summary

This is a brief status report on the results of Sandia's Red Team assessments of various DARPA ITO research projects. At this time, the project that raises the most questions is the MIT Lincoln Labs Intrusion Detection Evaluation project. Some of our basic concerns are listed here, but a more detailed assessment will be forthcoming. Other assessments are proceeding but are hampered by a lack of detailed design documentation.

Background

In early FY98, DARPA's ITO office contracted with Sandia National Laboratories to perform Red Team Quick Looks of various projects in ITO's collection of funded projects. After some negotiation, ITO and Sandia agreed to evaluate five projects. The particular projects were chosen because of their perceived impact on other DARPA programs.

This report is an interim status report on the progress of this effort by Sandia's Red Team. The goals of this report include:

- Inform the sponsor of Sandia's progress to date.
- Inform the sponsor of Sandia's plans for the continuation of this work.
- Invite the sponsor to provide feedback on the work done to date.
- Invite the sponsor to offer new direction for the future work.

Current Status

Projects

The Sandia Cyber Red Team has been studying five select DARPA ITO projects including:

- MIT Lincoln Labs Intrusion Detection Evaluation
- Generic Software Wrappers
- Agile (includes FLUKE)
- Ensemble
- IDIP

Challenges

The biggest challenge thus far has been getting accurate documentation in a timely fashion from the various developers. It has been difficult identifying, contacting, and then getting the attention of the principal investigators for most of these developmental projects.

Highlights

MIT Lincoln Labs Intrusion Detection Evaluation

We have many concerns about the validity of this work. We do not dispute the need for training data for intrusion detection and response (IDR) systems. However, we have these particular concerns with MIT/LL's approach:

- The number and types of attacks that are currently represented are limited and are not representative of current real-world threat vectors.
- The current operating environment does not represent most current corporate network environments. The test network assumes operating systems and hardware that are more representative of a typical university-computing environment.
- The current test network incorporates a naïve model for how Microsoft NT networking would be deployed.
- Our best information is that no commercial IDR systems have been evaluated using the MIT/LL model. However, MIT/LL has not been forthcoming with information on which IDR systems have actually been evaluated.
- Background data appears to be generated using weak statistical models or single observations. We would prefer that this data be generated using more sophisticated statistical models that are validated using multiple independent observations.

Our plans are to develop a more detailed quick look assessment report. This report would include the relative strengths, weaknesses, and our recommendations for improving the current approach. Our plan is to share this with MIT/LL and present our assessment with MIT/LL's comments by the end of April 1999. Ideally, we would like to work directly with MIT/LL to improve their process in an effort to achieve better evaluation results.

In addition, we speculate that MIT/LL may not be able to evaluate IDR systems based solely on their response to a given set of training data. It may be prudent to expand the scope of their evaluation effort to examine the architecture and methodology of a given IDR system under consideration. We will propose some additional evaluation processes and their potential benefits in our April report.

Generic Software Wrappers

This evaluation has gone well. The developers (Lee Badger & Mark Feldman of TIS Labs at Network Associates) have been forthcoming with a variety of documents on this project, including their own list of relative weaknesses of the technology.

At this point, wrappers appear to be a viable technique for protecting a few critical processes on a given computing platform. However, it appears to be far from being a generic security approach that can be widely adopted on a variety of platforms.

Our next steps here are to design some experiments to demonstrate the ease (or difficulty) of subverting wrappers in a typical workstation environment. We would like to run these experiments in Q4FY99.

Agile

The biggest challenge in assessing the Agile Security Architecture has been obtaining relevant technical documentation. Although the team at the University of Utah has responded to some of our inquires, none of the previously-supplied documents contain any relevant details that we are seeking.

We have finally been able to establish contact with the Principal Investigator, Jay Lepreau. Our current plan is to visit with Jay at his facility to determine if there is any more substantive documentation.

Eventually, we plan to stand up an Agile system. Since the project implements a POSIX subset in their operating system, we theorize that all of the traditional Unix attacks may be viable, including subverting file protections, gaining privileges, and other traditional attacks. We will know more once we have a system up and running.

Ensemble

Ensemble is turning out to be a difficult technology to evaluate. So far, all of the information that we have is fairly high level, whereas most of the potential risks appear to involve details of the implementation. Ensemble can be configured to meet a variety of needs, and this inherent flexibility actually hinders detailed analysis.

Our current plan for Ensemble is to (a) stand up a version of this product in our laboratory and then (b) study an actual implementation. The implementation under consideration is the Quoin integration effort for the Navy's Hiper-D program. Our expectation is that there should be fewer ambiguities in a deployed system, and this will aid our assessment. However, it is not clear that the Quoin project will share the detailed development information that we will need to perform our analysis.

IDIP

The Intrusion Detection and Isolation Protocol (IDIP) is a relatively mature technology. IDIP is being deployed in several developmental systems, including the DARPA/ISO Information Assurance architecture. The Principal Investigator, Dan Schnackenberg, has been forthcoming with relevant documentation.

At this point, parts of IDIP seem to be well organized. Our current approach is to postulate attacks against IDIP, and then design experiments that support or refute our assertions. Some candidate attacks include:

- Attacks on the key management system. Since the protocol uses a shared-secret authentication technique, it would be interesting to see if an adversary can compromise the shared secret and then manipulate the IDIP subnet by masquerading as a member of the IDIP community.
- Denial of Service through network congestion. It is not clear how much overhead traffic IDIP actually generates. If it is significant, we might be able to maliciously manipulate IDIP either as an insider or an outsider into generating enough traffic to disrupt the network.
- Disruption of the Community. An adversary may gain some advantage if they can disrupt communications between certain members of a given IDIP community. There are also various protocols for communicating between members of different IDIP communities. These potential protocol failures must be examined.
- Traffic Analysis. IDIP might actually provide an adversary with valuable information on how to attack a network.

We are looking forward to sharing our early results with the IDIP developers in an effort to identity any potential improvements.

Summary

Sandia's Red Team continues to evaluate the five selected ITO projects. The project that has raised the most concern is the MIT/LL project. A more detailed quick look assessment of the MIT/LL project will be forthcoming. Other assessments continue, pending the availability of detailed documentation.

Contacts

Comments, questions, or concerns arising from this document should be addressed to the author:

Bradley J. Wood
Information Design Assurance Red Team (IDART) Program Manager
Sandia National Laboratories
PO Box 5800, M/S 0449
Albuquerque, NM 87185-0449

Phone: 505/845-8461
Email: bjwood@sandia.gov